DeepCorrect: Correcting DNN Models against Image Distortions

Tejas Borkar, *Student Member, IEEE,* and Lina Karam, *Fellow, IEEE*

Abstract—In recent years, the widespread use of deep neural networks (DNNs) has facilitated great improvements in performance for computer vision tasks like image classification and object recognition. In most realistic computer vision applications, an input image undergoes some form of image distortion such as blur and additive noise during image acquisition or transmission. Deep networks trained on pristine images perform poorly when tested on such distortions. In this paper, we evaluate the effect of image distortions like Gaussian blur and additive noise on the activations of pre-trained convolutional filters. We propose a metric to identify the most noise susceptible convolutional filters and rank them in order of the highest gain in classification accuracy upon correction. In our proposed approach called *DeepCorrect*, we apply small stacks of convolutional layers with *residual connections*, at the output of these ranked filters and train them to correct the worst distortion affected filter activations, whilst leaving the rest of the pre-trained filter outputs in the network unchanged. Performance results show that applying *DeepCorrect* models for common vision tasks like image classification (ImageNet), object recognition (Caltech-101, Caltech-256) and scene classification (SUN-397), significantly improves the robustness of DNNs against distorted images and outperforms other alternative approaches.

Index Terms—deep neural networks, image distortion, image classification, residual learning, image denoising, image deblurring.

I. INTRODUCTION

TODAY, state-of-the-art algorithms for computer vision tasks like image classification, object recognition and semantic segmentation employ some form of deep neural networks (DNNs). The ease of design for such networks, afforded by numerous open source deep learning libraries [1], [2], has established DNNs as the go-to solution for many computer vision applications. Even challenging computer vision tasks like image classification [3], [4], [5], [6] and object recognition [7], [8], [9], which were previously considered to be extremely difficult, have seen great improvements in their state-of-the-art results due to the use of DNNs. An important factor contributing to the success of such deep architectures in computer vision tasks is the availability of large scale annotated datasets [10], [11].

The visual quality of input images is an aspect very often overlooked while designing DNN based computer vision systems. In most realistic computer vision applications, an input image undergoes some form of image distortion including blur and additive noise during image acquisition, transmission or storage. However, most popular large scale

T. Borkar and L. Karam are with the Department of Electrical, Computer and Energy Engineering, Arizona State University, Tempe, AZ, 85281 USA e-mail: {tsborkar, karam}@asu.edu.

datasets do not have images with such artifacts. Dodge and Karam [12] showed that even though such image distortions do not represent adversarial samples for a DNN, they do cause a considerable degradation in classification performance. Fig. 1 shows the effect of image quality on the prediction performance of a DNN trained on high quality images devoid of distortions.

Testing distorted images with a pre-trained DNN model for AlexNet [6], we observe that adding even a small amount of distortion to the original image results in a misclassification, even though the added distortion does not hinder the human ability to classify the same images (Fig. 1). In the cases where the predicted label for a distorted image is correct, the prediction confidence drops significantly as the distortion severity increases. This indicates that features learnt from a dataset of high quality images are not invariant to image distortion or noise and cannot be directly used for applications where the quality of images is different than that of the training images. Some issues to consider while deploying DNNs in noise/distortion affected environments include the following. For a network trained on undistorted images, are all convolutional filters in the network equally susceptible to noise or blur in the input image? Are networks able to learn some filters that are invariant to input distortions, even when such distortions are absent from the training set? Is it possible to identify and rank the convolutional filters that are most susceptible to image distortions and recover the lost performance, by only correcting the outputs of such ranked filters?

In our proposed approach called *DeepCorrect*, we try to address these aforementioned questions through the following novel contributions:

- Evaluating the effect of common image distortions like Gaussian blur and Additive White Gaussian Noise (AWGN) on the outputs of pre-trained convolutional filters. We find that for every layer of convolutional filters in the DNN, certain filters are more susceptible to input distortions than others and that correcting the activations of these filters can help recover lost performance.
- Measuring the susceptibility of convolutional filters to input distortions and ranking filters in the order of highest susceptibility to input distortion.
- Correcting the worst distortion-affected filter activations by appending *correction units*, which are small blocks of stacked convolutional layers trained using a target-oriented loss, at the output of select filters, whilst leaving the rest of the pre-trained filter outputs in a DNN unchanged.

Original — blur σ = 2 — blur σ = 4 — blur σ = 6

goldfish 0.977 — **goldfish 0.873** — **goldfish 0.398** — **goldfish 0.100**

gila monster 0.234 — *rock python 0.034* — *velvet 0.016* — *velvet 0.010*

(a)

Original — AWGN σ = 20 — AWGN σ = 60 — AWGN σ = 100

goldfish 0.977 — **goldfish 0.254** — *starfish 0.054* — *starfish 0.027*

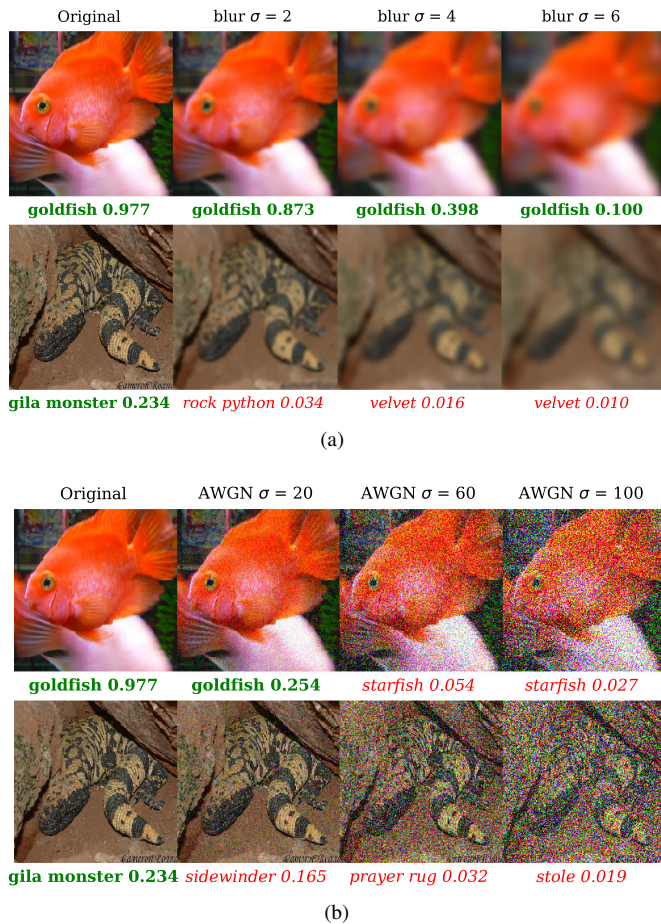**gila monster 0.234** — *sidewinder 0.165* — *prayer rug 0.032* — *stole 0.019*

(b)

Fig. 1. Effect of image quality on DNN predictions, with predicted label and confidence generated by a pre-trained AlexNet [6] model. Distortion severity increases from left to right, with the left-most image in a row having no distortion (original). Bold green text indicates correct classification, while italic red text denotes misclassification[1]. (a) Examples from the ILSVRC2012 [10] validation set distorted by Gaussian blur. (b) Examples from the ILSVRC2012 validation set distorted by Additive White Gaussian Noise (AWGN)

- Representing full-rank convolutions in our *DeepCorrect* models with rank-constrained approximations consisting of a sequence of separable convolutions with rectangular kernels to mitigate the additional computational cost introduced by our *correction units*. This results in pruned *DeepCorrect* models that have almost the same computational cost as the respective pre-trained DNN, during inference.

Applying our *DeepCorrect* models for common vision tasks like image classification [10], object recognition [13] [14] and scene classification [15] significantly improves the robustness of DNNs against distorted images and also outperforms other alternative approaches, while training significantly lesser parameters. To ensure reproducibility of presented results, the code for *DeepCorrect* is made publicly available at https://github.com/tsborkar/DeepCorrect.

The remainder of the paper is organized as follows. Section II provides an overview of the related work in assessing and improving the robustness of DNNs to input image perturbations. Section III describes the distortions, network architectures and datasets we use for analyzing the distortion susceptibility of convolutional filters in a DNN. A detailed

description of our proposed approach is presented in Section IV followed, in Section V, by extensive experimental validation with different DNN architectures and multiple datasets covering image classification, object recognition and scene classification. Concluding remarks are given in Section VI.

II. RELATED WORK

DNN susceptibility to specific small magnitude perturbations which are imperceptible to humans but cause networks to make erroneous predictions with high confidence (adversarial samples) has been studied in [16] and [17]. The concept of rubbish samples proposed in [18] studies the vulnerability of DNNs to make arbitrary high confidence predictions for random noise images that are completely unrecognizable to humans, i.e., the images contain random noise and no actual object. However, both adversarial samples and rubbish samples are relatively less encountered in common computer vision applications as compared to other common distortions due to image acquisition, storage, transmission and reproduction.

Karam and Zhu [19] present QLFW, a face matching dataset consisting of images with five types of quality distortions. Basu *et al.* [20] present the n-MNIST dataset, which adds Gaussian noise, motion blur and reduced contrast to the original images of the MNIST dataset. Dodge and Karam [12] evaluate the impact of a variety of quality distortions such as Gaussian blur, AWGN and JPEG compression on various state-of-the-art DNNs and report a substantial drop in classification accuracy on the ImageNet (ILSVRC2012) dataset in the presence of blur and noise. A similar evaluation for the task of face recognition is presented in [21].

Rodner *et al.* [22] assess the sensitivity of various DNNs to image distortions like translation, AWGN and salt & pepper noise, for the task of fine grained categorization on the CUB-200-2011 [23] and Oxford flowers [24] datasets, by proposing a first-order Taylor series based gradient approximation that measures the expected change in final layer outputs for small perturbations to input image. Since a gradient approximation assumes small perturbations, Rodner *et al.*'s sensitivity measure does not work well for higher levels of distortion as shown by [22] and does not assess susceptibility at a filter level within a DNN. Furthermore, Rodner *et al.* do not present a solution for making the network more robust to input distortions; instead, they simply fine-tune the whole network with the distorted images added as part of data augmentation during training. Retraining large networks such as VGG16 [3] or ResNet-50 [5] on large-scale datasets is computationally expensive. Unlike Rodner *et al.*'s work, our proposed ranking measure assesses sensitivity of individual convolutional filters in a DNN, is not limited to differentiable distortion processes, and holds good for both small and large perturbations.

Zheng *et al.* [25] improve the general robustness of DNNs to unseen small perturbations in the input image through the introduction of distortion agnostic *stability training*, which minimizes the KL-divergence between DNN output predictions for a clean image and a noise perturbed version of the same image. The perturbed images are generated by adding

---

[1]All figures in this paper are best viewed in color.

uncorrelated Gaussian noise to the original image. *Stability training* provides improved DNN robustness against JPEG compression, image scaling and cropping. However, the top-1 accuracy of *stability trained* models is lower than the original model, when tested on most distortions including motion blur, defocus blur and additive noise among others [26]. Sun *et al.* also propose a distortion agnostic approach to improve DNN robustness by introducing 3 feature quantization operations, i.e., a floor function with adaptive resolution, a power function with learnable exponents and a power function with data dependent exponents, that act on the convolutional filter activations before the ReLU non-linearity is applied. However, similar to *stability training*, the performance of *feature quantized* models is lower than the original model, when tested on distortions like defocus blur and additive noise. Additionally, no single feature quantization function consistently outperforms the other two for all types of distortion. Although both distortion agnostic methods [25], [26] improve DNN robustness, their top-1 accuracy is much lower than the accuracy of the original DNN on distortion free images, making it difficult to deploy these models.

A non-blind approach to improve the resilience of networks trained on high quality images would be to retrain the network parameters (fine-tune) on images with observed distortion types. Vasiljevic *et al.* [27] study the effect of various types of blur on the performance of DNNs and show that DNN performance for the task of classification and segmentation drops in the presence of blur. Vasiljevic *et al.* [27] and Zhou *et al.* [28] show that fine-tuning a DNN on a dataset comprised of both distorted and undistorted images helps to recover part of the lost performance when the degree of distortion is low.

Diamond *et al.* [29] propose a joint denoising, deblurring and classification pipeline. This involves an image preprocessing stage that denoises and deblurs the image in a manner that preserves image features optimal for classification rather than aesthetic appearance. The classification stage has to be fine-tuned using distorted and clean images, while the denoising and deblurring stages assume a priori knowledge of camera parameters and the blur kernel, which may not be available at the time of testing.

III. Experimental Setup

Here, we describe the various image distortions and network architectures used to evaluate the susceptibility of individual convolutional filters to input distortions. We use the ILSVRC-2012 validation set [10] for our experiments.

A. Distortions

We first focus on evaluating two important and conflicting types of image distortions: Gaussian blur and AWGN over 6 levels of distortion severity. Gaussian blur, often encountered during image acquisition and compression [30], represents a distortion that eliminates high frequency discriminative object features like edges and contours, whereas AWGN is commonly used to model additive noise encountered during image acquisition and transmission. We use a noise standard deviation $\sigma_n \in \{10, 20, 40, 60, 80, 100\}$ for AWGN and blur standard
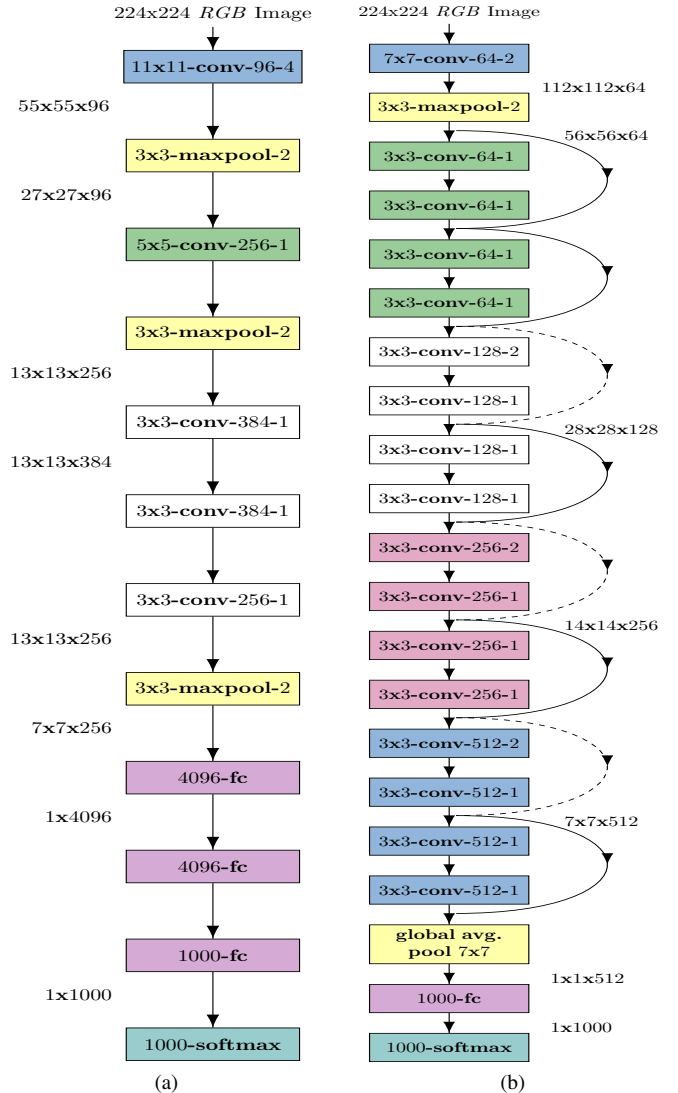

Fig. 2. Network architectures for our baseline models. Convolutional layers are parameterized by kxk-conv-d-s, where kxk is the spatial extent of the filter, d is the number of output filters in a layer and s represents the filter stride. Maxpooling layers are parameterized as kxk-maxpool-s, where s is the spatial stride. (a) AlexNet DNN [6][3] (b) ResNet18 [5][4].

deviation $\sigma_b \in \{1, 2, 3, 4, 5, 6\}$ for Gaussian blur. The size of the blur kernel is set to 4 times the value of σ_b.

B. Network Architectures

We use two different network architectures as shown in Fig. 2, specifically: a shallow 8-layered DNN (AlexNet)[2] and a deeper 18-layered network with residual connections (ResNet18) [5]. We use the term "pre-trained" or "baseline"

---

[2] We use the code and weights for the AlexNet DNN available online at https://github.com/heuritech/convnets-keras

[3] Every convolutional layer is followed by a ReLU non-linearity for the AlexNet DNN. In addition to the ReLU non-linearity, the first and second convolutional layers of the AlexNet DNN are also followed by a local response normalization operation [6].

[4] Every convolutional layer is followed by a batch normalization operation and ReLU non-linearity for the ResNet18 model. Skip connections and residual feature maps are combined through an element-wise addition. Dashed-line skip connections perform an identity mapping using a stride of 2 to reduce feature map size and pad zero entries along the channel axis to increase dimensionality [5].

TABLE I
TOP-1 ACCURACY OF PRE-TRAINED NETWORKS FOR DISTORTION
AFFECTED IMAGES AS WELL AS UNDISTORTED IMAGES (ORIGINAL). FOR
GAUSSIAN BLUR AND AWGN, ACCURACY IS REPORTED BY AVERAGING
OVER ALL LEVELS OF DISTORTION.

Models	Original	Gaussian blur	AWGN
AlexNet	0.5694	0.2305	0.2375
ResNet18	0.6912	0.3841	0.3255

network to refer to any network that is trained on undistorted images.

IV. DEEPCORRECT

Although pre-trained networks perform poorly on test images with significantly different image statistics than those used to train these networks (Table I), it is not obvious if only some convolutional filters in a network layer are responsible for most of the observed performance gap or if all convolutional filters in a layer contribute more or less equally to the performance degradation. If only a subset of the filters in a layer are responsible for most of the lost performance, we can focus on restoring only the most severely affected activations and avoid modifying all the remaining filter activations in a DNN.

A. Ranking Filters through Correction Priority

We define the output of a single convolutional filter $\phi_{i,j}$ to the input $\mathbf{x}_i$ by $\phi_{i,j}(\mathbf{x}_i)$, where i and j correspond to layer number and filter number, respectively. If $g_i(\cdot)$ is a transformation that models the distortion acting on filter input $\mathbf{x}_i$, then the output of a convolutional filter $\phi_{i,j}$ to the distortion affected input is given by $\widetilde{\phi_{i,j}}(\mathbf{x}_i) = \phi_{i,j}(g_i(\mathbf{x}_i))$. It should be noted that $\widetilde{\phi_{i,j}}(\mathbf{x}_i)$ represents the filter activations generated by distorted inputs and $\phi_{i,j}(\mathbf{x}_i)$ represents the filter activations for undistorted inputs. Assuming we have access to $\phi_{i,j}(\mathbf{x}_i)$ for a given set of input images, replacing $\widetilde{\phi_{i,j}}(\mathbf{x}_i)$ with $\phi_{i,j}(\mathbf{x}_i)$ in a deep network is akin to perfectly correcting the activations of the convolutional filter $\phi_{i,j}$ against input image distortions. Computing the output predictions by swapping a distortion affected filter output with its corresponding clean output for each of the ranked filters would improve classification performance. The extent of improvement in performance is indicative of the susceptibility of a particular convolutional filter to input distortion and its contribution to the associated performance degradation.

We now define the correction priority of a convolutional filter $\phi_{i,j}$ as the improvement in DNN performance on a validation set, generated by replacing $\widetilde{\phi_{i,j}}(\mathbf{x}_i)$ with $\phi_{i,j}(\mathbf{x}_i)$ for a pre-trained network. Let the baseline performance (computed over distorted images) for a network be p_b, which can be obtained by computing the average top-1 accuracy of the network over a set of images or another task-specific performance measure. Let $p_{swp}(i,j)$ denote the new improved performance of the network after swapping $\widetilde{\phi_{i,j}}(\mathbf{x}_i)$ with $\phi_{i,j}(\mathbf{x}_i)$. As our implementation focuses on classification tasks, the average top-1 accuracy over a set of distorted images is used to measure p_b and $p_{swp}(i,j)$. The correction priority for filter $\phi_{i,j}$ is then given by:

$$\tau(i,j) = p_{swp}(i,j) - p_b \qquad (1)$$

Algorithm 1 Computing Correction Priority

Input: $(\mathbf{x}_{1,i}, g_i(\mathbf{x_{1,i}}), y_1), \ldots, (\mathbf{x}_{M,i}, g_i(\mathbf{x_{M,i}}), y_M)$ are given triplets with $1 \leq i \leq L$, where i represents the layer number, $\mathbf{x_{m,i}}$ is the m^{th} undistorted input for layer i and $g_i(\mathbf{x_{m,i}})$ is the corresponding distorted version, M is the total number of images in the validation set and y_m is the ground-truth label for the m^{th} input image.
Output: Correction priority τ
1: $p_b := 0$
2: **for** $m = 1$ to M **do**
3:     Predict class label $ypred_m$ for distorted image $g_1(\mathbf{x}_{m,1})$
4:     Compute $p_b = p_b + \frac{1}{M}h(y_m, ypred_m)$,
5: where $h(y_m, ypred_m) = 1$, if $y_m = ypred_m$ and 0 otherwise.
6: **end for**
7: **for** $i = 1$ to L **do**
8:     $N_i \leftarrow$ number of filters in layer i
9:     $\phi_{i,j} \leftarrow j^{th}$ convolutional filter in the i^{th} layer
10:     **for** $j = 1$ to N_i **do**
11:         $p_{swp}(j) = 0$
12:         **for** $m = 1$ to M **do**
13:             $\phi_{i,j}(g_i(\mathbf{x}_{m,i})) \leftarrow \phi_{i,j}(\mathbf{x}_{m,i})$
14:             Predict class label $ypred_m$
15:             $p_{swp}(j) = p_{swp}(j) + \frac{1}{M}h(y_m, ypred_m)$,
16: where $h(y_m, ypred_m) = 1$, if $y_m = ypred_m$ and 0 otherwise.
17:         **end for**
18:     **end for**
19:     $\tau(i,j) \leftarrow p_{swp}(j) - p_b$
20: **end for**
21: **return** τ

A higher $\tau(i,j)$ indicates higher susceptibility of the convolutional filter $\phi_{i,j}$ to input distortion. Using the proposed ranking measure in Equation (1) and 5000 images (i.e., 5 images per class) randomly sampled from the ILSVRC-2012 training set, we compute correction priorities for every convolutional filter in the network and rank the filters in descending order of correction priority. The detailed overview and pseudo-code for computing correction priorities is summarized in Algorithm 1.

We evaluate in Fig. 3, the effect of correcting different percentages, β_i, of the ranked filter activations in the i^{th} DNN layer of AlexNet for distortion affected images. For the AlexNet model, it is possible to recover a significant amount of the lost performance, by correcting only 50% of the filter activations in any one layer, which indicates that a select subset of convolutional filters in each layer are indeed more susceptible to distortions than the rest. Although we show graphs for the AlexNet model, form our experiments, we make similar observations for the ResNet18 model as well.

Convolutional filter visualizations from the first layer of the pre-trained AlexNet model (Fig. 4a) reveal two types of filter kernels: 1) mostly color agnostic, frequency- and orientation-selective filters that capture edges and object contours and 2) color specific blob shaped filters that are sensitive to specific color combinations. Figs. 4b and 4c visualize the top 50% filters in the first convolutional layer of AlexNet, that are most susceptible to Gaussian blur and AWGN, respectively, as identified by our proposed ranking metric. The identified filters most susceptible to Gaussian blur are mainly frequency- and orientation-selective filters, most of which are color agnostic, while filters most susceptible to AWGN are a mix of both color specific blobs and frequency- and orientation-selective filters.

Fig. 3. Effect of varying the percentage of corrected filter activations $\beta_i \in \{10\%, 25\%, 50\%, 75\%, 90\%\}$, in the i^{th} convolutional layer (conv i) of pre-trained AlexNet, for AWGN and Gaussian blur affected images, respectively.


(a)

(b)

(c)

Fig. 4. (a) 96 convolutional filter kernels of size 11x11x3 in the first convolutional layer of pre-trained AlexNet. (b) Convolutional filter kernels most susceptible to Gaussian blur (top 50%), as identified by our proposed ranking metric. (c) Convolutional filter kernels most susceptible to AWGN (top 50%), as identified by our proposed ranking metric. In (b) and (c), the filters are sorted in descending order of susceptibility going row-wise from top left to bottom right.

This is in line with our intuitive understanding that Gaussian blur majorly affects edges and object contours and not object color, while AWGN affects color as well as object contours.

B. Correcting Ranked Filter Outputs

Here, we propose a novel approach, which we refer to as *DeepCorrect*, where we learn a task-driven corrective transform that acts as a distortion masker for convolutional filters

that are most susceptible to input distortion, while leaving all the other pre-trained filter outputs in the layer unchanged. Let R_i represent a set consisting of the N_i ranked filter indices in the i^{th} layer of the network, computed using the procedure in Section IV-A. Also let R_{i,β_i} represent a subset of R_i consisting of the top $\beta_i N_i$ ranked filter indices in network layer i, where N_i is the total number of convolutional filters in layer i and β_i is the percentage of filters corrected in layer i, as defined in Section IV-A. If Φ_i represents the set of convolutional filters in the i^{th} layer, the objective is to learn a transform $F_{corr_i}(:)$ such that:

$$F_{corr_i}(\Phi_{R_{i,\beta_i}}(g_i(\mathbf{x}_i))) \approx \Phi_{R_{i,\beta_i}}(\mathbf{x}_i) \qquad (2)$$

where $\mathbf{x}_i$ is the undistorted input to the i^{th} layer of convolutional filters and $g_i(\cdot)$ is a transformation that models the distortion acting on $\mathbf{x}_i$. Since we do not assume any specific form for the image distortion process, we let the corrective transform $F_{corr_i}(:)$ take the form of a shallow *residual block*, which is a small stack of convolutional layers (4 layers) with a single skip connection [5], such as the one shown in Fig. 5. We refer to such a *residual block* as a *correction unit*. $F_{corr_i}(:)$ can now be estimated using a target-oriented loss such as the one used to train the original network, through backpropogation [31], but with much less number of parameters.

Consider an L layered DNN Φ that has been pre-trained for an image classification task using clean images. Φ can be interpreted as a function that maps network input $\mathbf{x}$ to an output vector $\Phi(\mathbf{x}) \in \mathbb{R}^d$, such that:

$$\Phi = \Phi_L \circ \Phi_{L-1} \circ \ldots \Phi_2 \circ \Phi_1 \qquad (3)$$

where Φ_i is the mapping function (set of convolutional filters) representing the i^{th} DNN layer and d is the dimensionality of the network output.

Fig. 5. *Correction unit* based on a residual function [5], acting on the outputs of $\beta_i N_i$ $(0 < \beta_i < 1)$ filters out of N_i total filters in the i^{th} convolutional layer of a pre-trained DNN. All convolutional layers in the *residual block*, except the first and last layer, are parameterized by kxk-conv-D_i-s, where kxk is spatial extent of the filter, D_i (correction unit kernel depth) is the number of output filters in a layer, s represents the filter stride and i represents the layer number of the convolutional layer being corrected in the pre-trained DNN.

Without loss of generality, if we add a *correction unit* that acts on the top $\beta_1 N_1$ ranked filters in the first network layer, then the resultant network Φ_{corr} is given by:

$$\Phi_{corr} = \Phi_L \circ \Phi_{L-1} \circ \ldots \Phi_2 \circ \Phi_{1_{corr}} \tag{4}$$

where $\Phi_{1_{corr}}$ represents the new mapping function for the first layer, in which the corrective transform $F_{corr_1}(:)$ acts on the activations of the filter subset $\Phi_{R_{1,\beta_1}}$ and all the remaining filter activations are left unchanged. If $\mathbf{W}_1$ represents the trainable parameters in F_{corr_1}, then F_{corr_1} can be estimated by minimizing :

$$E(\mathbf{W_1}) = \lambda \mathscr{R}(\mathbf{W_1}) + \frac{1}{M} \sum_{m=1}^{M} \mathscr{L}(y_m, \Phi_{corr}(\mathbf{x}_m)) \tag{5}$$

where λ is a constant, $\mathscr{R}$ is a regularizer such as l_1 norm or l_2 norm, $\mathscr{L}$ is a standard cross-entropy classification loss, y_m is the target output label for the m^{th} input image $\mathbf{x}_m$, M represents the total number of images in the training set and, since we train on a collection of both distorted and clean images, $\mathbf{x}_m$ represents a clean or a distorted image. The trainable parameters in Equation (5) are $\mathbf{W}_1$, while all other network parameters are fixed and kept the same as those in the pre-trained models. Although Equation (5) shows *correction unit* estimation for only the first layer, it is possible to add such units at the output of distortion susceptible filters in any layer and in one or more layers. Fig. 6 shows an example of a *DeepCorrect* model for the pre-trained AlexNet model in Fig. 2a.

[5]The *correction unit* is applied before the ReLU non-linearity acting upon the distortion-susceptible convolutional filter outputs. Max pooling layers following convolutional layers 1, 2 and 5 in the pre-trained AlexNet have not been shown in the ImageNet *DeepCorrect* model for uniformity.
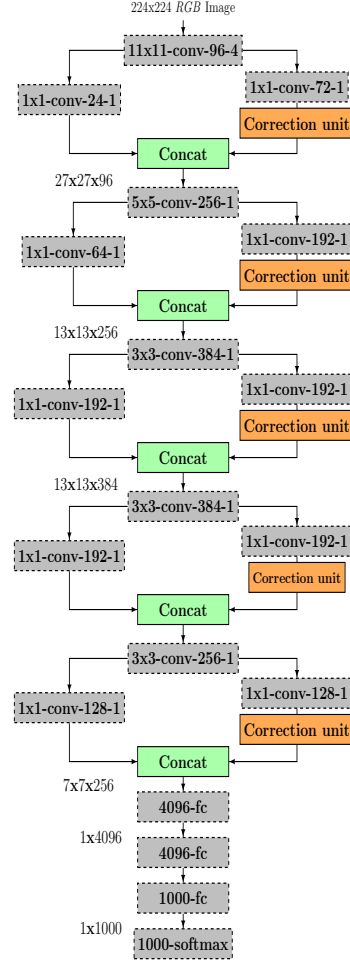

Fig. 6. *DeepCorrect* model for AlexNet, with 75% filter outputs corrected in the first two layers and 50% filter outputs corrected in the next three layers[5]. Convolution layers from the original architecture in Fig. 2, shown in gray with dashed outlines, are non-trainable layers and their weights are kept the same as those of the pre-trained model.

C. Rank-constrained DeepCorrect models

The inference time of our *DeepCorrect* models (Section IV-B) can be slower than the respective baseline models due to the additional computational cost introduced by the *correction units*. To mitigate the impact of these additional computations, we propose the use of a rank-constrained approximation to the full-rank *DeepCorrect* model, which not only has the same computational cost as the corresponding baseline DNN but also retains almost 99% of the performance of our full-rank *DeepCorrect* models.

Consider the n^{th} full-rank 3-D convolutional filter with weights $\mathbf{W}_n \in \mathbb{R}^{k \times k \times C}$ in a DNN convolutional layer with N filters, where $k \times k$ represents the filter's spatial extent and C is the number of input channels for the filter; then a rank constrained convolution is implemented by factorizing the convolution of $\mathbf{W}_n$ with input z into a sequence of separable convolutions (i.e., horizontal and vertical filters) as in [32]:

$$\mathbf{W}_n * z \approx \sum_{p=1}^{P} h_n^p * (v_p * z) = \sum_{p=1}^{P} h_n^p * \sum_{c=1}^{C} v_p^c * z^c \tag{6}$$

where the first convolutional filter bank consists of P vertical filters $\{v_p \in \mathbb{R}^{k \times 1 \times C} : p \in [1...P]\}$ and the second convolution

consists of a horizontal filter that operates on P input feature maps $\{h_n \in \mathbb{R}^{1 \times k \times P}\}$. The number of intermediate filter maps, P, controls the rank of the low-rank approximation. The computational cost of the original full-rank convolutional layer for N output feature maps with width W' and height H' is $O(Nk^2CH'W')$, whereas the rank-constrained approximation has a computational cost of $O((N+C)kPH'W')$ and a speedup can be achieved when $NkC > (N+C)P$. For the special case of convolutional layers in our *correction units*, where $N = C$ (Fig. 5), if $P = N/2$, the computational cost of a convolutional layer can be reduced k times (typically, $k = 3$).

Our rank-constrained *DeepCorrect* model is thus generated by replacing each full-rank convolutional layer (except 1×1 layers) with its respective rank-constrained approximation with P for each approximation chosen such that the total computational cost of our model is the same as the baseline DNN. Instead of using iterative methods or training the separable filters from random weights, we use the simple yet fast, matrix decomposition approach of Tai *et. al.* [33] to get the exact global optimizer of the rank-constrained approximation from its respective trained full-rank *DeepCorrect* model.

V. EXPERIMENTAL RESULTS

We evaluate *DeepCorrect* models against the alternative approaches of network fine-tuning and *stability training* [25], for the DNN architectures mentioned in Section III. The DNN models are trained and tested using a single Nvidia Titan-X GPU. Unlike common image denoising and deblurring methods like BM3D [34] and NCSR [35] which expect the distortion level to be known during both train and test phases or learning-based methods that train separate models for each distortion level, *DeepCorrect* trains a single model for all distortion levels at once and, consequently, there is no need to know the distortion level at test time.

A. AlexNet Analysis

1) Finetune model: We fine-tune the AlexNet model in Fig. 2a on a mix of distortion affected images and clean images to generate a single fine-tuned model and refer to this model as Finetune in our results. Starting with an initial learning rate (= 0.001) that is 10 times lower than that used to generate the pre-trained model in Fig. 2a a fixed number of iterations (62500 iterations $\approx$ 10 epochs), with the learning rate reduced by a factor of 10 after every 18750 iterations (roughly 3 epochs). We use a data augmentation method as proposed by [5].

2) Stability trained model: Following the *stability training* method outlined in [25] that considers that unseen distortions can be modelled by adding AWGN to the input image, we fine-tune all the fully connected layers of the pre-trained AlexNet model by minimizing the KL-divergence between the classification scores for a pair of images (I, I'), where $I' = I + \eta$ and $\eta \sim \mathcal{N}(0, \sigma^2)$. We use the same hyper-parameters used for the classification task in [25]: $\sigma^2 = 0.04$ and regularization coefficient $\alpha = 0.01$.

TABLE II
TOP-1 ACCURACY OF ALEXNET-BASED DNN MODELS FOR DISTORTION AFFECTED IMAGES OF THE IMAGENET VALIDATION SET (ILSVRC2012), AVERAGED OVER ALL LEVELS OF DISTORTION AND CLEAN IMAGES. BOLD NUMBERS SHOW BEST ACCURACY AND UNDERLINED NUMBERS SHOW NEXT BEST ACCURACY.

Method	Gaussian blur	AWGN
Baseline	0.2305	0.2375
Finetune	0.4596	0.4894
Finetune-rc	0.4549	0.4821
Deepcorr	**0.5071**	**0.5092**
Deepcorr-b	<u>0.5022</u>	<u>0.5063</u>
Deepcorr-rc	0.4992	0.5052
Stability [25]	0.2163	0.2305
NCSR [35]+AlexNet [6]	0.2193	-
BM3D [34]+AlexNet [6]	-	0.5032

TABLE III
COMPUTATIONAL PERFORMANCE OF ALEXNET-BASED DNN MODELS.

Metric	Baseline/ Finetune	*Deepcorr*	*Deepcorr*-b	*Deepcorr*-rc
FLOPs	7.4×10^8	23.9×10^8	11.8×10^8	7.8×10^8
Trainable params	60.96M	2.81M	1.03M	1.03M [6]

3) DeepCorrect models: Our main *DeepCorrect* model for AlexNet shown in Fig. IV-B and referred to as *Deepcorr* in Table II is generated by correcting 75% ranked filter outputs in the first two layers ($\beta_1, \beta_2 = 0.75$) and 50% ranked filter outputs in the next three layers ($\beta_3, \beta_4, \beta_5 = 0.5$) of the pre-trained AlexNet shown in Fig. 2a. The *correction units* (Fig. 5) in each convolutional layer are trained using an initial learning rate of 0.1 and the same learning rate schedule, data augmentation and total iterations used for generating the Finetune model in Section V-A1. We also generate two additional variants based on our *Deepcorr* model, *Deepcorr*-b, a computationally lighter model than *Deepcorr* based on a bottleneck architecture for its *correction units* as described later in this section, and *Deepcorr*-rc, which is a rank-constrained model (Section IV-C) derived from the full-rank *Deepcorr*-b model such that its test-time computational cost is almost the same as the Finetune model. For comparison, a rank-constrained model is also derived for the Finetune model using similar decomposition parameters as *Deepcorr*-rc and the resulting model is denoted by Finetune-rc.

Table II shows the superior performance of our proposed method as compared to the alternative approaches and Table III summarizes the computational performance in terms of trainable parameters and floating point operations (FLOPs) of these DNN models during training and testing, respectively. We evaluate the training computational cost in terms of the number of trainable parameters that are updated during training. The test-time computational cost is evaluated in terms of the total FLOPs, i.e., total multiply-add operations needed for a single image inference, as outlined by He *et. al.* in [5]. In particular, a $k \times k$ convolutional layer operating on C input maps and producing N output feature maps of width W' and height H'

---

[6]Since *Deepcorr*-rc is derived from the *Deepcorr*-b model through a low-rank approximation, the number of trainable parameters and subsequently its effect on model convergence is the same as *Deepcorr*-b.

TABLE IV

CORRECTION UNIT ARCHITECTURES FOR ALEXNET-BASED DEEPCORRECT MODELS. THE NUMBER FOLLOWING CORR-UNIT SPECIFIES THE LAYER AT WHICH CORRECTION UNITS ARE APPLIED. CORRECTION UNIT CONVOLUTIONAL LAYERS ARE REPRESENTED AS kxk, d, WHERE kxk IS THE SPATIAL EXTENT OF A FILTER AND d IS THE NUMBER OF FILTERS IN A LAYER. STACKED CONVOLUTIONAL LAYERS ARE ENCLOSED IN BRACKETS, FOLLOWED BY THE NUMBER OF LAYERS STACKED.

Model	Corr-unit 1 $\beta_1 N_1 = 72$ Output size 55×55	Corr-unit 2 $\beta_2 N_2 = 192$ Output size 27×27	Corr-unit 3 $\beta_3 N_3 = 192$ Output size 13×13	Corr-unit 4 $\beta_4 N_4 = 192$ Output size 13×13	Corr-unit 5 $\beta_5 N_5 = 128$ Output size 13×13
Deepcorr Total FLOPs: 16.5×10^8	$1 \times 1, 72$ $\left[5 \times 5, 72\right] \times 2$ $1 \times 1, 72$ FLOPs: 8.1×10^8	$1 \times 1, 192$ $\left[3 \times 3, 192\right] \times 2$ $1 \times 1, 192$ FLOPs: 5.3×10^8	$1 \times 1, 192$ $\left[3 \times 3, 192\right] \times 2$ $1 \times 1, 192$ FLOPs: 1.2×10^8	$1 \times 1, 192$ $\left[3 \times 3, 192\right] \times 2$ $1 \times 1, 192$ FLOPs: 1.2×10^8	$1 \times 1, 128$ $\left[3 \times 3, 128\right] \times 2$ $1 \times 1, 128$ FLOPs: 5.5×10^7
Deepcorr-b Total FLOPs: 4.4×10^8	$1 \times 1, 36$ $\left[3 \times 3, 36\right] \times 3$ $1 \times 1, 72$ FLOPs: 1.2×10^8	$1 \times 1, 96$ $\left[3 \times 3, 96\right] \times 3$ $1 \times 1, 192$ FLOPs: 2.0×10^8	$1 \times 1, 96$ $\left[3 \times 3, 96\right] \times 3$ $1 \times 1, 192$ FLOPs: 4.8×10^7	$1 \times 1, 96$ $\left[3 \times 3, 96\right] \times 3$ $1 \times 1, 192$ FLOPs: 4.8×10^7	$1 \times 1, 64$ $\left[3 \times 3, 64\right] \times 3$ $1 \times 1, 128$ FLOPs: 2.1×10^7
Deepcorr-rc Total FLOPs: 2.5×10^8	$1 \times 1, 36$ $\left[\begin{smallmatrix}3 \times 1, 27 \\ 1 \times 3, 36\end{smallmatrix}\right] \times 3$ $1 \times 1, 72$ FLOPs: 6.8×10^7	$1 \times 1, 96$ $\left[\begin{smallmatrix}3 \times 1, 72 \\ 1 \times 3, 96\end{smallmatrix}\right] \times 3$ $1 \times 1, 192$ FLOPs: 1.1×10^8	$1 \times 1, 96$ $\left[\begin{smallmatrix}3 \times 1, 72 \\ 1 \times 3, 96\end{smallmatrix}\right] \times 3$ $1 \times 1, 192$ FLOPs: 2.7×10^7	$1 \times 1, 96$ $\left[\begin{smallmatrix}3 \times 1, 72 \\ 1 \times 3, 96\end{smallmatrix}\right] \times 3$ $1 \times 1, 192$ FLOPs: 2.7×10^7	$1 \times 1, 64$ $\left[\begin{smallmatrix}3 \times 1, 48 \\ 1 \times 3, 64\end{smallmatrix}\right] \times 3$ $1 \times 1, 128$ FLOPs: 1.2×10^7

requires $Nk^2 CH'W'$ FLOPs [5]. The detailed architecture and corresponding FLOPs for each *correction unit* in our different *DeepCorrect* models for AlexNet are summarized in Table IV. Design choices for various *correction unit* architectures and their impact on inference-time computational cost are also discussed later in this section.

As it can be seen from Table II, the *Deepcorr* model, which is our best performing model in terms of top-1 classification accuracy, outperforms the Finetune model with $\approx 10\%$ and $\approx 4\%$ relative average improvement for Gaussian blur and AWGN, respectively, by just training ≈ 2.81M parameters (Table III) as compared to 61M parameters for the Finetune model (i.e., 95.4% lesser parameters). *Deepcorr* significantly improves the robustness of a pre-trained DNN achieving an average top-1 accuracy of 0.5071 and 0.5092 for Gaussian blur and AWGN affected images, respectively, as compared to the corresponding top-1 accuracy of 0.2305 and 0.2375 for the pre-trained AlexNet DNN.

All our *DeepCorrect* model variants consistently outperform the Finetune and Stability models for both distortion types (Table II). We also observe that fine-tuning DNN models on distortion specific data significantly outperforms DNN models trained through distortion agnostic *stability training*. For completeness, we also compare classification performance with AlexNet when combined with a commonly used non-blind image denoising method (BM3D) proposed by [34] and a deblurring method (NCSR) proposed by [35], where BM3D and NCSR are applied prior to the baseline AlexNet, for AWGN and Gaussian blur, respectively. Table II shows top-1 accuracy for these two methods with the *Deepcorr* model outperforming each for AWGN and Gaussian blur, respectively.

4) Correction unit architectures: Increasing the correction unit kernel depth (D_i in Fig. 5) makes our proposed *correction unit* fatter, whereas decreasing D_i makes the *correction unit* thinner. A natural choice for D_i would be to make it equal to the number of distortion susceptible filters that need correction ($\beta_i N_i$), in a DNN layer i; this is also the default parameter setting used in the *correction units* for *Deepcorr*. Since D_i is always equal to the number of distortion susceptible filters that need correction (i.e., $D_i = \beta_i N_i$), the number of trainable parameters in the *correction units* of *Deepcorr* scale linearly with the number of corrected filters ($\beta_i N_i$) in each convolutional layer, even though *Deepcorr* trains significantly lesser parameters than Finetune and still achieves a better classification accuracy (Tables II and III).

As shown in Table III, during the testing phase, the *correction units* in *Deepcorr* add almost 2 times more FLOPs relative to the baseline DNN, for evaluating a single image. As discussed in more detail later in this section, one way to limit the number of trainable parameters and FLOPs is to explore a bottleneck architecture for our *correction units*, where the convolutional kernel depth D_i (Fig. 5) is set to 50% of the distortion susceptible filters that need correction in a DNN layer i (i.e., $D_i = \frac{\beta_i N_i}{2}$) as compared to D_i being set to the number of filters to correct (i.e., $D_i = \beta_i N_i$) in *Deepcorr*. Replacing each *correction unit* in *Deepcorr* with such a bottleneck *correction unit*, which we refer to as *Deepcorr*-b, results in a *DeepCorrect* model that has significantly less trainable parameters and FLOPs than the original *Deepcorr* model (Table IV). Compared to the *correction units* in *Deepcorr*, bottleneck *correction units* provide a 60% reduction in FLOPs on average, with a 85% reduction in FLOPs for Corr-unit 1 as shown in Table IV. From Tables II and III, it can be seen that *Deepcorr*-b achieves $\approx 99\%$ of the average accuracy achieved by *Deepcorr*, with a 63% reduction in trainable parameters and a 73% reduction in FLOPs, relative to *Deepcorr*. As shown in Table III, the *Deepcorr*-b model still requires 58% more FLOPs relative to the baseline DNN, for evaluating a single image at test time. A further reduction in the computational cost at test-time can be achieved by deriving a rank-constrained *DeepCorrect* model (*Deepcorr*-rc) from the *Deepcorr*-b model using the approach outlined in Section IV-C. By replacing each full rank convolutional layer in a *Deepcorr*-b *correction unit* with a pair of separable convolutions (Section IV-C), we can reduce the FLOPs for each *correction unit* by an additional 46% as shown in Table IV. Replacing the pre-trained convolutional

Fig. 7. Effect of *DeepCorrect* ranking metric on *correction unit* performance when integerated with AlexNet [6]. Dashed lines represent *correction units* trained on the least susceptible filters in a DNN layer and solid lines represent *correction units* trained on the most susceptible filters, as identified by our ranking metric (Section IV-A).

layers of the baseline AlexNet (which are left unchanged in the *DeepCorrect* models and shown in gray in Fig. 6) by their equivalent low-rank approximations provides an additional 29% reduction in FLOPs, relative to the baseline AlexNet model, such that the resultant *Deepcorr*-rc model now has almost the same FLOPs as Finetune (Table III) and still retains almost 99% of the accuracy achieved by *Deepcorr*.

5) Effect of ranking on correction unit performance: If the superior performance of our *DeepCorrect* model is only due to the additional network parameters provided by the *correction unit*, then re-training the correction unit on the least susceptible $\beta_i N_i$ filter outputs in a DNN layer should also achieve the same performance as that achieved by applying a *correction unit* on the $\beta_i N_i$ most susceptible filter outputs identified by our ranking metric (Section IV-A). β_i and N_i, as defined in Section IV-A, represent the percentage of filters corrected in the i^{th} layer and the total number of filters in the i^{th} layer, respectively. To this end, we evaluate the performance of training *correction units* on: 1) $\beta_i N_i$ filters most susceptible to distortion, and 2) $\beta_i N_i$ filters least susceptible to distortion, as identified by our ranking metric. For this analysis, $\beta_i \in \{50\%, 75\%\}$.

As shown in Fig. 7, *correction units* trained on the $\beta_i N_i$ most distortion susceptible filters (solid lines) outperform those trained on the least susceptible filters (dashed lines) of conv 1, conv 2 and conv 5 layers of the AlexNet model, for Gaussian blur and AWGN, respectively. Although we observe a similar trend for the conv 3 and conv 4 layers, we plot results for only conv 1, conv 2 and conv 5 as these show the largest difference in performance due to ranking distortion susceptible filters. Correcting the top 75% distortion susceptible filters (solid orange), as identified by our ranking measure, achieves the best performance in all layers. Similarly, for all layers and distortions, *correction units* trained on the top 50% susceptible
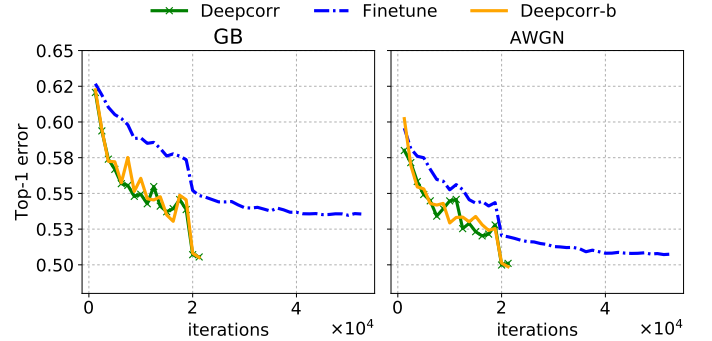

Fig. 8. Top-1 error on the ILSVRC-2012 validation set, for *DeepCorrect* model variants and fine-tuning.

filters (solid green) not only significantly outperform *correction units* trained on the 50% least susceptible filters (dashed blue) but also outperform *correction units* trained on the 75% least susceptible filters (dashed red), where 25% filters are shared among both approaches.

6) Accelerating training: We analyze the evolution of the validation set error over training iterations for *Deepcorr*, *Deepcorr*-b as well as the Finetune model and stop training for the *Deepcorr* models when their respective validation error is lesser than or equal to the minimum validation error achieved by the Finetune model. For any particular value of validation error achieved by the Finetune model in Fig. 8, both the *Deepcorr* model variants are able to achieve the same validation error in much lesser number of training iterations. Thus, we conclude that just learning corrective transforms for activations of a subset of convolutional filters in a layer accelerates training through reduced number of training epochs needed for convergence.

7) Generalization of DeepCorrect Features to other Datasets: To analyze the ability of distortion invariant features

TABLE V
MEAN ACCURACY PER CATEGORY OF PRE-TRAINED ALEXNET DEEP
FEATURE EXTRACTOR FOR CLEAN IMAGES.

Caltech-101	Caltech-256	SUN-397
0.8500	0.6200	0.3100

TABLE VI
MEAN ACCURACY PER CATEGORY FOR GAUSSIAN BLUR AFFECTED
IMAGES, AVERAGED OVER ALL DISTORTION LEVELS. BOLD NUMBERS
SHOW BEST ACCURACY.

Dataset	Baseline	Finetune	*Deepcorr*
Caltech-101	0.4980	0.7710	**0.8371**
Caltech-256	0.2971	0.5167	**0.5883**
SUN-397	0.1393	0.2369	**0.3049**

TABLE VII
MEAN ACCURACY PER CATEGORY FOR AWGN AFFECTED IMAGES,
AVERAGED OVER ALL DISTORTION LEVELS. BOLD NUMBERS SHOW BEST
ACCURACY.

Dataset	Baseline	Finetune	*Deepcorr*
Caltech-101	0.3423	0.7705	**0.8034**
Caltech-256	0.1756	0.4995	**0.5482**
SUN-397	0.0859	0.1617	**0.2936**

TABLE VIII
TOP-1 ACCURACY OF RESNET18-BASED DNN MODELS FOR DISTORTION
AFFECTED IMAGES OF THE IMAGENET VALIDATION SET (ILSVRC2012),
AVERAGED OVER ALL LEVELS OF DISTORTION AND CLEAN IMAGES. BOLD
NUMBERS SHOW BEST ACCURACY AND UNDERLINED NUMBERS SHOW
NEXT BEST ACCURACY.

Method	**G.Blur**	**AWGN**	**M.Blur**	**D.Blur**	**Cam.Blur**
Baseline	0.3841	0.3255	0.4436	0.3582	0.4749
Finetune	0.5617	0.5970	0.6197	0.5615	0.6041
Finetune-rc	0.5548	0.5898	0.6113	0.5537	0.5973
Deepcorr	0.5808	<u>0.6058</u>	**0.6498**	**0.6005**	<u>0.6346</u>
Deepcorr-b	**0.5839**	**0.6087**	<u>0.6474</u>	<u>0.5831</u>	**0.6365**
Deepcorr-rc	<u>0.5821</u>	0.6033	0.6411	0.5785	0.6276
Stability [25]	0.3412	0.3454	0.4265	0.3182	0.4720

TABLE IX
COMPUTATIONAL PERFORMANCE OF RESNET18-BASED DNN MODELS.

Metric	Baseline/ Finetune	*Deepcorr*	*Deepcorr*-b	*Deepcorr*-rc
FLOPs	1.8×10^9	3.5×10^9	2.9×10^9	1.8×10^9
Trainable params	11.7M	8.41M	5.5M	5.5M

learnt for image classification to generalize to related tasks like object recognition and scene recognition, we evaluate the *Deepcorr* and Finetune models trained on the ImageNet dataset (Section V-A) as discriminative deep feature extractors on the Caltech-101 [13], Caltech-256 [14] and SUN-397 [15] datasets. Unlike object recognition datasets like Caltech-101 and Caltech-256, which bear some similarity to an image classification/object recognition dataset like ImageNet, a scene recognition dataset like SUN-397 bears no similarity to the ImageNet dataset and is expected to be challenging for features extracted using models learnt on ImageNet [36]. Following the experimental procedure proposed by [36], we use the output of the first (Caltech-101 and Caltech-256) or second (SUN-397) fully-connected layer in these models as a deep feature extractor for images affected by distortion.

Since the above deep feature models have not been trained on any one of these datasets (Caltech-101, Caltech-256 and SUN-397), for each dataset, we train linear SVMs on top of the deep features, which are extracted from a random set of training data, and evaluate the performance in terms of mean accuracy per category averaged over 5 data splits, following the training procedure adopted by [36]. The training data for each split consists of 25 training images and 5 validation images per class, sampled randomly from the considered dataset and all remaining images are used for testing. A baseline accuracy for undistorted images is first established by training linear SVMs on features extracted only from undistorted images using the AlexNet DNN shown in Fig. 2a, and results are reported in Table V. Similar to the evaluation in Section V-A, we now independently add Gaussian blur and AWGN to train and test images using the same distortion levels as reported in Section III and report performance averaged over all distortion levels in Tables VI-VII for deep features extracted using baseline AlexNet, Finetune and *Deepcorr* models trained on ImageNet[7]. Both

---

[7]For each of the three models, a single set of linear SVMs is trained for images affected by different levels of distortion and also clean images.

Gaussian blur and AWGN significantly affect the accuracy of the baseline feature extractor for all three datasets, with a 41% and 60% drop in respective accuracies for Caltech-101, a 52% and 71% drop in respective accuracies for Caltech-256, and a 55% and 72% drop in respective mean accuracy for SUN-397, relative to the benchmark performance for clean images. For Caltech-101, the *Deepcorr* feature extractor outperforms the Finetune feature extractor with a 8.5% and 4.2% relative improvement in mean accuracy for Gaussian blur and AWGN affected images, respectively. For Caltech-256, the *Deepcorr* feature extractor outperforms the Finetune feature extractor with a 13.8% and 9.7% relative improvement for Gaussian blur and AWGN, respectively. Similarly, features extracted using the *Deepcorr* model significantly outperform those extracted using the Finetune model for SUN-397, with a 28.7% and 81.5% relative improvement in mean accuracy for Gaussian blur and AWGN, respectively. The large performance gap between *Deepcorr* and Finetune feature extractors highlights the generic nature of distortion invariant features learnt by our *DeepCorrect* models.

B. ResNet18 Analysis

Similar to the AlexNet analysis in Section V-A, for ResNet18, we evaluate the performance of our proposed *DeepCorrect* models against DNN models trained through fine-tuning and *stability training*. In addition to Gaussian blur (GB) and AWGN, we also evaluate 3 additional distortion types used by Vasiljevic *et. al.* [27], namely 1) Motion blur (M.Blur), 2) Defocus blur (D.Blur) and 3) Camera shake blur (Cam.Blur). Spatially uniform disk kernels of varying radii are used to simulate defocus blur while horizontal and vertical box kernels of single pixel width and varying length are used to simulate motion blur (uniform linear motion) [27]. For simulating camera shake blur, we use the code in [37] to generate 100 random blur kernels and report performance averaged over all camera shake blur kernels.

Using the training procedures outlined in Sections V-A1 and V-A2, we generate fine-tuned and *stability trained* models from the baseline ResNet18 DNN by training all layers (11.7M

TABLE X

CORRECTION UNIT ARCHITECTURES FOR RESNET18-BASED DEEPCORRECT MODELS. THE NUMBER FOLLOWING CORR-UNIT SPECIFIES THE LAYER AT WHICH CORRECTION UNITS ARE APPLIED. CORRECTION UNIT CONVOLUTIONAL LAYERS ARE REPRESENTED AS kxk, d, WHERE kxk IS THE SPATIAL EXTENT OF A FILTER AND d IS THE NUMBER OF FILTERS IN A LAYER. STACKED CONVOLUTIONAL LAYERS ARE ENCLOSED IN BRACKETS, FOLLOWED BY THE NUMBER OF LAYERS STACKED.

Model	Corr-unit 1 $\beta_1 N_1 = 48$ Output size 112×112	Corr-unit 3, 5 $\beta_2 N_2 = 48$ Output size 56×56	Corr-unit 5, 7 $\beta_3 N_3 = 96$ Output size 28×28	Corr-unit 9, 11 $\beta_4 N_4 = 192$ Output size 14×14	Corr-unit 13, 15 $\beta_5 N_5 = 384$ Output size 7×7
Deepcorr Total FLOPs: 1.7×10^9	$1 \times 1, 48$ $\begin{bmatrix} 3 \times 3, 48 \end{bmatrix} \times 2$ $1 \times 1, 48$ FLOPs: 5.7×10^8	$1 \times 1, 48$ $\begin{bmatrix} 3 \times 3, 48 \end{bmatrix} \times 2$ $1 \times 1, 48$ FLOPs: 1.4×10^8	$1 \times 1, 96$ $\begin{bmatrix} 3 \times 3, 96 \end{bmatrix} \times 2$ $1 \times 1, 96$ FLOPs: 1.4×10^8	$1 \times 1, 192$ $\begin{bmatrix} 3 \times 3, 192 \end{bmatrix} \times 2$ $1 \times 1, 192$ FLOPs: 1.4×10^8	$1 \times 1, 384$ $\begin{bmatrix} 3 \times 3, 384 \end{bmatrix} \times 2$ $1 \times 1, 384$ FLOPs: 1.4×10^8
Deepcorr-b Total FLOPs: 1.1×10^9	$1 \times 1, 31$ $\begin{bmatrix} 3 \times 3, 31 \end{bmatrix} \times 3$ $1 \times 1, 48$ FLOPs: 3.6×10^8	$1 \times 1, 31$ $\begin{bmatrix} 3 \times 3, 31 \end{bmatrix} \times 3$ $1 \times 1, 48$ FLOPs: 9.2×10^7	$1 \times 1, 62$ $\begin{bmatrix} 3 \times 3, 62 \end{bmatrix} \times 3$ $1 \times 1, 96$ FLOPs: 9.2×10^7	$1 \times 1, 124$ $\begin{bmatrix} 3 \times 3, 124 \end{bmatrix} \times 3$ $1 \times 1, 192$ FLOPs: 9.2×10^7	$1 \times 1, 248$ $\begin{bmatrix} 3 \times 3, 248 \end{bmatrix} \times 3$ $1 \times 1, 384$ FLOPs: 9.2×10^7
Deepcorr-rc Total FLOPs: 0.6×10^9	$1 \times 1, 31$ $\begin{bmatrix} 3 \times 1, 24 \\ 1 \times 3, 31 \end{bmatrix} \times 3$ $1 \times 1, 48$ FLOPs: 2.0×10^8	$1 \times 1, 31$ $\begin{bmatrix} 3 \times 1, 24 \\ 1 \times 3, 31 \end{bmatrix} \times 3$ $1 \times 1, 48$ FLOPs: 5.1×10^7	$1 \times 1, 62$ $\begin{bmatrix} 3 \times 1, 48 \\ 1 \times 3, 62 \end{bmatrix} \times 3$ $1 \times 1, 96$ FLOPs: 5.1×10^7	$1 \times 1, 124$ $\begin{bmatrix} 3 \times 1, 96 \\ 1 \times 3, 124 \end{bmatrix} \times 3$ $1 \times 1, 192$ FLOPs: 5.1×10^7	$1 \times 1, 248$ $\begin{bmatrix} 3 \times 1, 192 \\ 1 \times 3, 248 \end{bmatrix} \times 3$ $1 \times 1, 384$ FLOPs: 5.1×10^7

parameters) of the DNN on a mix of distorted and clean images. Similarly, using the training procedure of Section V-A3, our competing *DeepCorrect* model is generated by training *correction units* that are appended at the output of the most susceptible filters in the odd-numbered convolutional layers (1 to 17) of the baseline ResNet18 model right after each skip connection merge (Fig. 2b). Similar to the AlexNet analysis presented earlier (Section V-A), we generate three *DeepCorrect models* (i.e., *Deepcorr*, *Deepcorr*-b and *Deepcorr*-rc). For comparison, a rank-constrained model is derived for the Finetune model using similar decomposition parameters as *Deepcorr*-rc and the resulting model is denoted by Finetune-rc. Table VIII summarizes the accuracy of ResNet18-based DNN models against various distortions, while Table IX shows the computational performance of the same DNN models, measured in terms of FLOPs and trainable parameters. Similar to the AlexNet analysis, the detailed architecture and corresponding FLOPs for each *correction unit* in our different *DeepCorrect* models for ResNet18 are shown in Table X.

From Table VIII, we observe that our *DeepCorrect* models not only significantly improve the robustness of the baseline DNN model to input distortions but also outperform the alternative approaches of model fine-tuning and *stability training* in terms of classification accuracy. *Deepcorr*, which trains almost 28% lesser parameters (Table IX) than Finetune, outperforms Finetune by 3.2% for Gaussian blur, 1.47% for AWGN, 4.86% for motion blur, 6.95% for defocus blur and 5% for camera shake blur. *Deepcorr*-b, which trains 50% lesser parameters (Table IX) than Finetune, outperforms Finetune by 3.95% for Gaussian blur, 1.9% for AWGN, 4.46% for motion blur, 3.84% for defocus blur and 5.36% for camera shake blur. Similarly our rank constrained model (*Deepcorr*-rc), which trains 50% parameters less than Finetune (Table IX), outperforms Finetune by 3.63% for Gaussian blur, 1% for AWGN, 3.45% for motion blur, 3% for defocus blur and 3.89% for camera shake blur. As shown in Tables IX and X, during the testing phase, *Deepcorr*-b requires 60% more FLOPs relative to the baseline DNN, for evaluating a single image. On the other hand, for *Deepcorr*-rc, just using rank-constrained *correction units* results in a 50% reduction in FLOPs as compared to *Deepcorr*-b (Table X), with a 30% additional reduction in FLOPs relative to baseline ResNet18 achieved by replacing the pre-trained convolutional layers of baseline ResNet18 with their low-rank approximations. From Tables VIII and IX, it can be seen that *Deepcorr*-rc requires almost the same number of FLOPs as Finetune but achieves a superior performance than the Finetune and Stability models without sacrificing inference speed.

VI. CONCLUSION

Deep networks trained on pristine images perform poorly when tested on distorted images affected by image blur or additive noise. Evaluating the effect of Gaussian blur and AWGN on the activations of convolutional filters trained on undistorted images, we observe that select filters in each DNN convolutional layer are more susceptible to input distortions than the rest. We propose a novel objective metric to assess the susceptibility of convolutional filters to distortion and use this metric to identify the filters that maximize DNN robustness to input distortions, upon correction of their activations.

We design *correction units*, which are *residual blocks* that are comprised of a small stack of trainable convolutional layers and a single skip connection per stack. These *correction units* are added at the output of the most distortion susceptible filters in each convolutional layer, whilst leaving the rest of the pre-trained (on undistorted images) filter outputs in the network unchanged. The resultant DNN models which we refer to as *DeepCorrect* models, significantly improve the robustness of DNNs against image distortions and also outperform the alternative approach of network fine-tuning on common vision tasks like image classification, object recognition and scene classification, whilst training significantly less parameters and achieving a faster convergence in training. Fine-tuning limits the ability of the network to learn invariance to severe levels of distortion, and re-training an entire network can be computationally expensive for very deep networks. By correcting the most distortion-susceptible convolutional filter outputs, we are not only able to make a DNN robust to severe distortions, but

are also able to maintain a very good performance on clean images.

Although we focus on image classification and object recognition, our proposed approach is not only generic enough to apply to a wide selection of tasks that use DNN models such as object detection [7], [38] or semantic segmentation [11] but also other less commonly occurring noise types like adversarial noise.

REFERENCES

[1] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, "Caffe: Convolutional architecture for fast feature embedding," *arXiv preprint arXiv:1408.5093*, 2014.

[2] F. Chollet, "Keras," https://github.com/fchollet/keras, 2015.

[3] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," CoRR, vol. abs/1409.1556, 2014. [Online]. Available: http://arxiv.org/abs/1409.1556

[4] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 1–9.

[5] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.

[6] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.

[7] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2014, pp. 580–587.

[8] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," in Advances in Neural Information Processing Systems, 2015, pp. 91–99.

[9] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 779–788.

[10] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. IEEE, 2009, pp. 248–255.

[11] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: Common objects in context," in European Conference on Computer Vision. Springer, 2014, pp. 740–755.

[12] S. Dodge and L. Karam, "Understanding how image quality affects deep neural networks," in Eighth International Conference on Quality of Multimedia Experience (QoMEX), 2016, pp. 1–6.

[13] L. Fei-Fei, R. Fergus, and P. Perona, "Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories," in Computer Vision and Image Understanding, vol. 106, no. 1. Elsevier, 2007, pp. 59–70. [Online]. Available: http://www.vision.caltech.edu/Image_Datasets/Caltech101/

[14] G. Griffin, A. Holub, and P. Perona, "Caltech-256 object category dataset," 2007. [Online]. Available: http://www.vision.caltech.edu/Image_Datasets/Caltech256/

[15] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba, "Sun database: Large-scale scene recognition from abbey to zoo," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, 2010, pp. 3485–3492.

[16] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," *arXiv preprint arXiv:1312.6199*, 2013.

[17] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," *arXiv preprint arXiv:1412.6572*, 2014.

[18] A. Nguyen, J. Yosinski, and J. Clune, "Deep neural networks are easily fooled: High confidence predictions for unrecognizable images," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2015, pp. 427–436.

[19] L. J. Karam and T. Zhu, "Quality labeled faces in the wild (QLFW): a database for studying face recognition in real-world environments," in Proc. SPIE 9394, Human Vision and Electronic Imaging, vol. XX, 2015.

[20] S. Basu, M. Karki, S. Ganguly, R. DiBiano, S. Mukhopadhyay, S. Gayaka, R. Kannan, and R. Nemani, "Learning sparse feature representations using probabilistic quadtrees and deep belief nets," in Neural Processing Letters. Springer, 2015, pp. 1–13.

[21] S. Karahan, M. K. Yildirum, K. Kirtac, F. S. Rende, G. Butun, and H. K. Ekenel, "How image degradations affect deep CNN-based face recognition?" in International Conference of the Biometrics Special Interest Group (BIOSIG). IEEE, 2016, pp. 1–5.

[22] E. Rodner, M. Simon, R. B. Fisher, and J. Denzler, "Fine-grained recognition in the noisy wild: Sensitivity analysis of convolutional neural networks approaches," *arXiv preprint arXiv:1610.06756*, 2016.

[23] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie, "The caltech-ucsd birds-200-2011 dataset," *Tech. Rep.*, 2011. [Online]. Available: http://www.vision.caltech.edu/visipedia/CUB-200.html

[24] M.-E. Nilsback and A. Zisserman, "Automated flower classification over a large number of classes," in Proceedings of the Indian Conference on Computer Vision, Graphics and Image Processing, Dec 2008, pp. 722–729.

[25] S. Zheng, Y. Song, T. Leung, and I. J. Goodfellow, "Improving the robustness of deep neural networks via stability training," *CoRR*, vol. abs/1604.04326, 2016. [Online]. Available: http://arxiv.org/abs/1604.04326

[26] Z. Sun, M. Ozay, Y. Zhang, X. Liu, and T. Okatani, "Feature quantization for defending against distortion of images," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.

[27] I. Vasiljevic, A. Chakrabarti, and G. Shakhnarovich, "Examining the impact of blur on recognition by convolutional networks," *CoRR*, vol. abs/1611.05760, 2016. [Online]. Available: http://arxiv.org/abs/1611.05760

[28] Y. Zhou, S. Song, and N.-M. Cheung, "On classification of distorted images with deep convolutional neural networks," *arXiv preprint arXiv:1701.01924*, 2017.

[29] S. Diamond, V. Sitzmann, S. P. Boyd, G. Wetzstein, and F. Heide, "Dirty pixels: Optimizing image classification architectures for raw sensor data," *CoRR*, vol. abs/1701.06487, 2017. [Online]. Available: http://arxiv.org/abs/1701.06487

[30] N. Ponomarenko, V. Lukin, A. Zelensky, K. Egiazarian, M. Carli, and F. Battisti, "TID2008-A database for evaluation of full-reference visual quality assessment metrics," in Advances of Modern Radioelectronics, vol. 10, no. 4, 2009, pp. 30–45.

[31] D. E. Rumelhart, R. Durbin, R. Golden, and Y. Chauvin, "Backpropagation," Y. Chauvin and D. E. Rumelhart, Eds. Hillsdale, NJ, USA: L. Erlbaum Associates Inc., 1995, ch. Backpropagation: The Basic Theory, pp. 1–34.

[32] M. Jaderberg, A. Vedaldi, and A. Zisserman, "Speeding up convolutional neural networks with low rank expansions," *CoRR*, vol. abs/1405.3866, 2014. [Online]. Available: http://arxiv.org/abs/1405.3866

[33] C. Tai, T. Xiao, X. Wang, and W. E, "Convolutional neural networks with low-rank regularization," *CoRR*, vol. abs/1511.06067, 2015. [Online]. Available: http://arxiv.org/abs/1511.06067

[34] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "BM3D image denoising with shape-adaptive principal component analysis," in Proc. Workshop on Signal Processing with Adaptive Sparse Structured Representations (SPARS09), 2009.

[35] W. Dong, L. Zhang, G. Shi, and X. Li, "Nonlocally centralized sparse representation for image restoration," in IEEE Trans. on Image Processing, vol. 22, no. 4, 2013, pp. 1620–1630.

[36] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell, "DeCAF: A deep convolutional activation feature for generic visual recognition," in International Conference on Machine Learning, 2014, pp. 647–655.

[37] A. Chakrabarti, "A neural approach to blind motion deblurring," *CoRR*, vol. abs/1603.04771, 2016. [Online]. Available: http://arxiv.org/abs/1603.04771

[38] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The Pascal visual object classes (VOC) challenge," in International Journal of Computer Vision, vol. 88, no. 2, Jun. 2010, pp. 303–338.